

WHAT IS CLAIMED IS:

1. A system, comprising:

5 a node including an active device, an interface to an inter-node network, a memory, and an address network coupling the active device, the interface, and the memory;

an additional node coupled to the node by the inter-node network;

10 wherein in response to receiving from the active device an address packet initiating a transaction to gain an access right to a coherency unit, the memory is configured to send data corresponding to the coherency unit to the active device dependent on memory response information associated 15 with the coherency unit;

wherein if the transaction cannot be satisfied within the node, the memory is configured to forward a report corresponding to the address packet to the interface, wherein in response to the report, the interface is configured to send the additional node a coherency message requesting the access right 20 via the inter-node network.

2. The system of claim 1, wherein the node includes a data network coupling the active device, the interface, and the memory, and wherein the memory is configured to 25 send the report to the interface in a data packet.

3. The system of claim 1, wherein the memory response information is configured to identify one of at least two possible response states.

4. The system of claim 1, wherein if the memory response information indicates a no response state, the memory is configured to send neither the data nor the report in response to the address packet;

5 wherein when the memory response information indicates the no response state, an additional active device included in the node has an ownership responsibility for the coherency unit and is configured to supply the data to the active device in response to receiving the address packet.

10 5. The system of claim 4, wherein if the address packet is conveyed in point-to-point mode by the address network, the memory is configured to send an additional address packet indicating the transaction to the additional active device having the ownership responsibility for the coherency unit.

15 6. The system of claim 4, wherein the memory is configured to update the memory response information to indicate the no response state in response to receiving the address packet, and wherein the active device is configured to gain an ownership responsibility for the address packet in response to receiving the address packet.

20 7. The system of claim 1, wherein the memory is configured to update the memory response information to indicate a response state in response to receiving an address packet providing the memory with an ownership responsibility for the coherency unit;

25 wherein if the memory response information identifies the response state, the memory is configured to send the data corresponding to the coherency unit to the active device in response to the address packet.

30 8. The system of claim 1, wherein if the memory response information indicates a shared response state, the memory is configured to not send the report and to send the data corresponding to the coherency unit if the access right is a read access right;

wherein if the memory response information indicates the shared response state and the access right is a write access right, the memory is configured to not send the data to the active device and to send the report to the interface.

5

9. The system of claim 8, wherein the memory is configured to update the response information to indicate the shared response state in response to a proxy address packet sent by the interface indicating that an additional active device in the additional node is requesting read access to the coherency unit.

10

10. The system of claim 1, wherein if the memory response information indicates an invalid response state, the memory is configured to not send the data to the active device and to send the report to the interface in response to the address packet.

15

11. The system of claim 10, wherein the memory is configured to update the memory response information to indicate the invalid response state in response to receiving a proxy address packet from the interface indicating that an additional active device included in the additional node is requesting write access to the coherency unit.

20

12. The system of claim 1, wherein the memory is configured to send the data dependent on both the memory response information and a global access state of the coherency unit within the node, wherein the memory is configured to send the report dependent on the global access state.

25

13. The system of claim 1, wherein the memory is configured to include a value of the memory response information in the report, wherein the value is a value of the memory response information before the memory response information is modified in response to the address packet.

14. The system of claim 13, wherein the interface is configured to maintain a plurality of records in an outstanding transaction queue, wherein each of the plurality of records corresponds to a respective report received from the memory, and wherein the interface is configured to determine a global access state of the coherency unit in the node from the 5 value of the memory response information included in the report and from one or more of the records in the outstanding transaction queue.

15. The system of claim 14, wherein in response to receiving an additional coherency message specifying the coherency unit from the additional node, the interface is configured to select a type of proxy address packet to send on the address network dependent on the global access state. 10

16. The system of claim 13, wherein if the memory subsequently updates the memory response information again in response to another address packet, the memory is configured to provide a new value of the memory response information to the interface. 15

17. A node for use in a multi-node system, the node comprising:

20 a plurality of devices including a memory, an active device, and an interface to an inter-node network coupling nodes in the multi-node computer system;

an address network configured to convey packets between the plurality of devices;

wherein in response to receiving from the active device an address packet 25 initiating a transaction to gain an access right to a coherency unit, the memory is configured to send data corresponding to the coherency unit to the active device dependent on memory response information associated with the coherency unit;

5

wherein if the transaction cannot be satisfied within the node, the memory is configured to forward a report corresponding to the address packet to the interface, wherein in response to the report, the interface is configured to send a coherency message requesting the access right to another one of the nodes via the inter-node network.

10

18. The node of claim 17, further comprising a data network coupling the active device, the interface, and the memory, wherein the memory is configured to send the report to the interface in a data packet.

20

19. The node of claim 17, wherein the memory response information is configured to identify one of at least two possible response states.

20

20. The node of claim 17, wherein if the memory response information indicates a no response state, the memory is configured to send neither the data nor the report in response to the address packet;

wherein when the memory response information indicates the no response state, an additional active device included in the node has an ownership responsibility for the coherency unit and is configured to supply the data to the active device in response to receiving the address packet.

25

21. The node of claim 20, wherein if the address packet is conveyed in point-to-point mode by the address network, the memory is configured to send an additional address packet indicating the transaction to the additional active device having the ownership responsibility for the coherency unit.

22. The node of claim 20, wherein the memory is configured to update the memory response information to indicate the no response state in response to receiving the address

packet, and wherein the active device is configured to gain an ownership responsibility for the address packet in response to receiving the address packet.

23. The node of claim 17, wherein the memory is configured to update the memory
5 response information to indicate a response state in response to receiving an address
packet providing the memory with an ownership responsibility for the coherency unit;

10 wherein if the memory response information identifies the response state, the
memory is configured to send the data corresponding to the coherency unit
to the active device in response to the address packet.

24. The node of claim 17, wherein if the memory response information indicates a
shared response state, the memory is configured to send the data corresponding to the
coherency unit if the access right is a read access right;

15 wherein if the memory response information indicates the shared response state
and the access right is a write access right, the memory is configured to
send the report to the interface.

20 25. The node of claim 24, wherein the memory is configured to update the response
information to indicate the shared response state in response to a proxy address packet
sent by the interface indicating that an additional active device in another node is
requesting read access to the coherency unit.

25 26. The node of claim 17, wherein if the memory response information indicates an
invalid response state, the memory is configured to send the report to the interface in
response to the address packet.

27. The node of claim 26, wherein the memory is configured to update the memory
30 response information to indicate the invalid response state in response to receiving a

proxy address packet from the interface indicating that an additional active device included in another node is requesting write access to the coherency unit.

28. The node of claim 17, wherein the memory is configured to send the data
5 dependent on both the memory response information and a global access state of the coherency unit within the node.

29. The node of claim 17, wherein the memory is configured to include a value of the memory response information in the report, wherein the value is a value of the memory
10 response information before the memory response information is modified in response to the address packet.

30. The node of claim 29, wherein the interface is configured to maintain a plurality
of records in an outstanding transaction queue, wherein each of the plurality of records
15 corresponds to a respective report received from the memory, and wherein the interface is configured to determine a global access state of the coherency unit in the node from the value of the memory response information included in the report and from one or more of the records in the outstanding transaction queue.

20 31. The node of claim 30, wherein in response to receiving an additional coherency message specifying the coherency unit from another node, the interface is configured to select a type of proxy address packet to send on the address network dependent on the global access state.

25 32. The node of claim 14, wherein if the memory subsequently updates the memory response information again in response to another address packet, the memory is configured to provide a new value of the memory response information to the interface.

33. A method of operating a multi-node system, wherein the multi-node system includes a node and an additional node coupled by an inter-node network, the method comprising:

5 an active device included in the node initiating a transaction to gain an access right to a coherency unit by sending an address packet on an address network included in the node;

10 dependent on memory response information associated with the coherency unit, a memory included in the node sending data corresponding to the coherency unit to the active device in response to the address packet;

15 if the transaction cannot be satisfied within the node, the memory forwarding a report corresponding to the address packet to the interface;

in response to the report, the interface sending the additional node a coherency message requesting the access right via the inter-node network.

34. The method of claim 33, wherein said forwarding the report comprises the 20 memory send a data packet including the report to the interface via a data network included in the node.

35. The method of claim 33, wherein the memory response information identifies one of at least two possible response states.

25
36. The method of claim 33, wherein if the memory response information indicates a no response state, the memory sends neither the data nor the report in response to the address packet;

wherein when the memory response information indicates the no response state, an additional active device included in the node has an ownership responsibility for the coherency unit and supplies the data to the active device in response to receiving the address packet.

5

37. The method of claim 36, further comprising: the memory sending an additional address packet indicating the transaction to the additional active device having the ownership responsibility for the coherency unit if the address packet is conveyed in point-to-point mode by the address network.

10

38. The method of claim 36, further comprising:

the memory updating the memory response information to indicate the no response state in response to receiving the address packet; and

15

the active device gaining an ownership responsibility for the address packet in response to receiving the address packet.

39. The method of claim 33, further comprising:

20

the memory updating the memory response information to indicate a response state in response to receiving an address packet providing the memory with an ownership responsibility for the coherency unit;

25

wherein if the memory response information identifies the response state, the memory sends the data corresponding to the coherency unit to the active device in response to the address packet.

40. The method of claim 33, wherein if the memory response information indicates a shared response state, the memory does not send the report and sends the data corresponding to the coherency unit if the access right is a read access right;

5 wherein if the memory response information indicates the shared response state and the access right is a write access right, the memory does not send the data to the active device and sends the report to the interface.

10 41. The method of claim 40, further comprising the memory updating the response information to indicate the shared response state in response to a proxy address packet sent by the interface indicating that an additional active device in the additional node is requesting read access to the coherency unit.

15 42. The method of claim 33, wherein if the memory response information indicates an invalid response state, the memory does not send the data to the active device and sends the report to the interface in response to the address packet.

20 43. The method of claim 42, further comprising the memory updating the memory response information to indicate the invalid response state in response to receiving a proxy address packet from the interface indicating that an additional active device included in the additional node is requesting write access to the coherency unit.

25 44. The method of claim 33, wherein the memory sends the data dependent on both the memory response information and a global access state of the coherency unit within the node, wherein the memory sends the report dependent on the global access state.

30 45. The method of claim 33, further comprising the memory including a value of the memory response information in the report, wherein the value is a value of the memory response information before the memory response information is modified in response to the address packet.

46. The method of claim 45, further comprising:

the interface maintaining a plurality of records in an outstanding transaction
5 queue, wherein each of the plurality of records corresponds to a respective report received from the memory; and

10 the interface determining a global access state of the coherency unit in the node from the value of the memory response information included in the report and from one or more of the records in the outstanding transaction queue.

47. The method of claim 46, further comprising:

15 in response to receiving an additional coherency message specifying the coherency unit from the additional node, the interface selecting a type of proxy address packet to send on the address network dependent on the global access state.

48. The method of claim 46, further comprising:

20 if the memory subsequently updates the memory response information again in response to another address packet, the memory providing a new value of the memory response information to the interface.